Particle Physics Data Grid Collaboratory Pilot

Quarterly Status Report of the Steering Committee,

combined: October – December 2003 and January –March 2004

16th March 2004

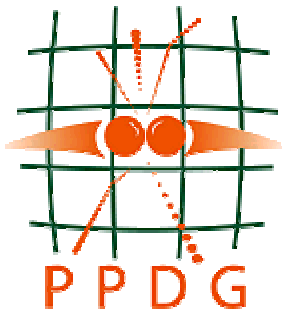

1 Project Overview

This report covers the accomplishments and activities of PPDG from October 2003 to the middle of March 2004. There was considerable progress in these two quarters as is shown in the milestone table in the next section. Several milestones culminated in the many demonstrations at the SC2003 conference in Phenix. All teams increased the level of their production activities over the Grid.

A significant concentration of effort in this period was the planning of future activity over the next 2.5 years to help achieve the reliable, persistent shared grid infrastructure that we envision as forming the backbone of distributed computing for high energy and nuclear physics. This planning is encapsulated in the renewal proposal[1] we submitted in January. A quick view of these plans is included in this report by including the milestones for PPDG years 4 and 5 in the Appendix. A significant new element to PPDG in these plans is to work on integrating the major laboratory computing facilities used by PPDG participants into a shared grid. This activity is called PPDG-OSG and is aligned with the Open Science Grid[2] roadmap.

News updates for the Virtual Data Toolkit and Grid2003 were presented to our sponsors.

1.1 Progress on the Year 3 Plan

Milestones identified in the Year 3 plan for this reporting period are shown with their status and comments:

Date	Team(s)	Task Name	Section + Milestone Number	Date Finished
8/03	Exec	Start monthly job management coordination meetings.	2.7.1	8/03
8/03	ATLAS & CMS	Grid2003 grid with four sites.	2.8.1	8/03
8/03	JLAB	SRM Improvements	4.4.1	[1], 11/03
8/18/03	ATLAS & CMS	Grid3 Integration Week.	2.8.2	8/22/03
9/03	Exec	Restart monthly monitoring coordination meetings.	2.7.2	[3]
9/03	Exec	Start monthly data management coordination meetings.	2.7.3	9/03
9/03	BaBar	Phased integration with the existing JImport tools replacing the pieces that perform crude SRB-like functions by SRB. With simulated p2p	4.2.2	[2], [4] 2/04
10/03	D0	D0 Finish JIM V1 installation	4.5.1	10/03
10/03	JLAB & STAR	Deliver requirements/definition for U-JDL batch job; gather PPDG feedback	2.9.1	9/03
10/03	SRB	Requirements for BaBar production	3.3.1	11/03
10/03	SRM	Complete the SRM v2.0 specification	3.4.1	10/03

[1] http://www.ppdg.net/docs/Proposals/scidac04_ppdg-final.pdf
[2] www.opensciencegrid.org

Date	Team(s)	Task Name	Section + Milestone Number	Date Finished
10/03 (10/04)	PPDG & DOESG	Document describing VO management practices	3.8.1	[5]
10/03	BaBar	BaBar Testing of new SRB with peer-to-peer support	4.2.4	11/03
10/03	JLAB	Replica Catalog Web Service Interface	4.4.2	[12]
10/21/03	ATLAS & CMS	Grid2003 SC2003 demonstrations.	2.8.3	11/03
11/03	Exec	Review progress of team plans and provide status summary.[6]	2.7.4	10/03
11/03	JLAB & STAR	JLAB-STAR Design implementation completed and first version of WSDL for batch web service	2.9.2	12/03
11/03 (4/04)	Condor	VDT based on GT3	3.5.1	[7]
11/03	D0	D0 using VDT	3.5.2	[11]
11/03	BaBar	BaBar using VDT	3.5.3	
11/03	Atlas	ATLAS Production system for simulation using Grid3	4.1.1	11/03
11/03	CMS	Working Grid3 application demonstrator for SC2003	4.3.1	11/03
12/03	Exec	Semi-annual PPDG collaboration/steering meeting.	2.7.5	12/03
12/03	Exec	Agree on project and team performance metrics	2.7.6	[8]
12/03	Exec	Agree on process for integration of middleware into the experiments' software stack.	2.7.7	1/04
12/03	JLAB & STAR	First implementations of site interface to map WSDL/XML onto PBS/LSF/Condor submissions	2.9.3	[9]
12/03	SRB	Initial implementation and deployment for BaBar production needs.	3.3.2	12/03
12/03	SRM	Developing a WSDL-based wrapper for DRM and HRM v1.0	3.4.2	12/03
12/03	PPDG & DOESG	DOESG-PPDG: MyProxy with credential repository	3.8.2	[10]
1/04	Exec	Publish PPDG documents on software testing, QA and middleware integration.	2.7.8	2/04
1/04	STAR	STAR Using VDT	3.5.4	12/03
1/04	JLab	Prototype LQCD meta data catalog	4.4.3	
1/04	D0	JIM V1.2 release	4.5.2	
1/04	STAR	Initial Replica Registration service (RRS) deployed	4.6.1	12/03

Comments

[1] from Chip: dependent on SRM v2.1….(please redefine to) 10/03 and we'll see if we can make it -- I expect the next spec version soon, and we would only need (I believe) to tweak our implementation to declare success.

[2] from Adil: ..milestone has slipped by 1.5 months due to my not taking into account ccin2p3 holiday schedules We have tested all the pieces, we just need to run them together. I'm not sure how prod it will be by Oct (maybe flaky

[3] from Ruth: Focus currently on Grid2003 and MonaLisa deployment and integration. Will hold meeting post-sc2003

[4] Wilko: JImport is a tool used at IN2P3 to import Objectivity databases. With the new computing model the BaBar events are stored in root files and there is no need to import Objectivity files anymore. Therefore we decided not to spent any effort on the integration of SRB and JImport.

[5] Doug: Deferred to Oct. 2004 as part of DOESG Phase2.

[6] In quarterly report, http://www.ppdg.net/docs/ppdg_qtrly_jul-sep-03.pdf

[7] VDT 1.3 planned for Spring 2004, see http://www.lsc-group.phys.uwm.edu/vdt/newsletter/2004-02.html

[8] First draft, Jan 2004, http://www.ppdg.net/docs/Papers/PPDG-36-23jan04.pdf

[9] JL: STAR will not be able to do this on time. JLab expects implementation available about May 1, 2004.

[10] Planned for 4/04 at NERSC.

[11] Deferred until after Mar 04 deployment for MC production.

[12] Deferred due to low priority.

1.2 Papers and Documents

The following reports and papers were posted at http://www.ppdg.net/docs/documents_and_information.htm.

Reports, Documents and Papers		Date/Version
PPDG-42	The SAM-Grid Fabric Services	pdf, ps
PPDG-41	Integration and Testing of Middleware into the Experiments' Software Stack	doc pdf
PPDG-40	PPDG Accomplishments - background for renewal proposal	doc pdf
PPDG-39	uJDL - STAR, JLAB	doc pdf
PPDG-38	Grid Analysis Environment Work Plan in association with the PPDG Proposal	doc pdf
PPDG-37	GRID2003 Lessons Learned	doc pdf steering meeting
PPDG-36	PPDG Project Performance Metrics	doc pdf Draft 23Jan04
PPDG-35	Interactive Analysis APIs Diagram and Description	pdf ppt
PPDG-34	Observations on Using Chimera with LCG-1 Sites Supporting ATLAS	pdf doc 20Jan04

2 Focus Areas and Common Projects

2.1 Data Management

In collaboration with the SRM SciDAC ISIC the SRM V2 specification was completed. http://sdm.lbl.gov/srm/documents/joint.docs/SRM.spec.v2.1.final.doc . Implementations at JLAB, Fermilab, LBNL and CERN are being upgraded to conform to the specification.

Progress was made on the Replica Registration Service and a prototype put into production for STAR.

U.S. ATLAS deployed the Globus Replica Location services as part of their Chimera based simulation distributed computing environment.

2.2 Job Management.

STAR coordinated the writing of a specification for a User Job Description Language (U-JDL) in collaboration with U.S. ATLAS and JLAB. This was presented and discussed at the PPDG collaboration meeting in December.

The STAR scheduler was integrated with Condor-G and demonstrated problems with the interface between GRAM and the LSF job scheduler in the presence of many jobs (the famous Globus Bugzilla 950). After much discussion and thinking STAR fixed a problem in the Globus-LSF adaptor – a contribution to the community. The STAR scheduled based on Condor-G is being testing. in progress, deployment expected Q1CY2004.

2.3 Production Grids

The Grid2003 project deployed a multi-VO application grid of 26 sites , 4 VOs and 2500 cpus over a common grid environment Grid3. http://www.ivdl.org/grid2003. The grid continues to be used for U.S. CMS simulation production, with opportunistic use of Grid3 resources.

The D0 SAMGrid JIM is deployed at 4 sites - Lancaster, Imperial College, GridKA and Wisconsin for monte carlo simulations. Performance and stability metrics are being measured. A wider deployment within the Dzero collaboration is planned over the next few months.

The Replica Registration Service prototype is deployed for STAR use between LBNL and BNL.

2.4 Data Analysis Working Group

The Data Analysis Working Group (CS-11) continued its useful series of meetings to provide a forum for discussion and planning across the experiment groups prototyping solutions to interactive analysis over grid environments. The two implementations of the Catalog Service have given an initial demonstration of the appropriateness and completeness of the interface developed in PPDG.

The JAS-Tech-X and Caltech Grid Analysis projects implemented the common Catalog Service the defined list of CS-11 services in PPDG-19.

2.5 Monitoring

MDS and MonaLisa were deployed and used for operational monitoring during Grid2003.

2.6 AAA

Progress continued outside of PPDG funded effort on necessary components of a production grid infrastructure. The joint US ATLAS, iVDGL, Fermilab, US CMS Virtual Organization registration project provided support for Grid3 http://www.uscms.org/s&c/VO/ ; DOE Science Grid is developing a Myproxy credential store at NERSC which may be used by STAR; Brookhaven continued towards deployment of the

GUMS software, and Fermilab put the Kerberos CA in production in addition to establishing an RA for DOEGrids host and service certificates.

3 Collaborations

3.1 DOE Science Grid

Collaboration with the DOE Science Grid continued. A phone meeting was held to discuss the services offered by the collaboratory and how requests and expectations of the user community can be made and met. Bimonthly phone meetings of Trillium and DOE Science Grid were started in January as a regular forum for discussion of mutual issues and plans.

3.2 Trillium and Grid2003

The Grid2003 project was successfully completed. The demonstrations at SC2003 were very successful. The following news items were posted.

An Application Grid Laboratory for Science

The Grid2003 project has deployed an international Data Grid with dozens of sites and thousands of processors. The U.S. Grid projects iVDGL, GriPhyN and PPDG, and the U.S participants in the LHC experiments ATLAS and CMS operate the facility jointly. Project highlights include:

- Participation by more than 25 sites across the US and Korea which collectively provide more than 2000 CPUs
- Resources used by 7 different scientific applications, including 3 high energy physics simulations and 4 data analyses in high energy physics, biochemistry, astrophysics and astronomy
- More than 100 individuals are currently registered with access to the Grid
- A peak throughput of 500-900 jobs running concurrently with a completion efficiency of approximately 75%

Cross Posted News Item between GRID2003 and LCG

Scientists from the LHC experiments and computer scientists from the Condor Project and the Globus Alliance have worked together towards demonstrations at the SIS-forum at CERN in Geneva, Switzerland. The demonstrations show interoperability between the U.S. based multi-organization Grid2003 shared grid demonstrator and grid services for the LHC Computing Grid Project (LCG). They use common grid middleware packaged through the Virtual Data Toolkit, built on the NSF Middleware Initiative release. Each grid relies on components developed by the other, as well as developments by joint projects across the groups. LCG also uses components developed in the European DataGrid project.

The demonstrations show initial small steps towards interoperating grid systems for the LHC experiments. One demonstration shows high speed, managed data movement by the CMS experiment between CERN and Fermilab in Illinois, using a common storage management interface, with an implementation developed by DESY, in Germany, and Fermilab connecting to a local implementation for the mass storage system at CERN. A second demonstration by the ATLAS experiment shows progress towards submitting physics simulation jobs across LCG and Grid2003 using additional Virtual Data services developed by the U.S. physics grid projects (GriPhyN, iVDGL and PPDG).

US-CMS has recently deployed a Storage Element service demonstrator

This Storage Element (SE) has several components: a Disk Management system, dCache, developed as a joint project with DESY and FNAL and Storage Resource Manager (SRM) as the common interface. dCache allows many storage pools to be grouped together into a storage system. The pools can be everything from high performance RAID devices to individual disks on worker nodes. dCache handles load balancing by replicating files between pools if they are frequently requested. It also transparently handles the loss of a pool, provided the data is available on another pool or a MSS. dCache was designed as a front

end to tape systems, which is the implementation at FNAL. The SRM interface is becoming the standard storage interface protocol. The implementation deployed in the US handles advanced functionality of proxy driven transfers, resource reservation, and transfer queuing. The functionality offered by this storage element solves many problems, which we've seen with the CMS application. Applications can reserve space for a period of time, ensuring that space will be available as the application runs. Worker nodes inside a NAT or with no external network connection at all can trigger data be replicated to the storage element, and then access the data through the LAN. Moving the storage element functionality off the compute element also reduces the load and improves the reliability of the node.

CMS has demonstrated interoperability between CERN and FNAL, through the SRM interface data was transferred from tape at CERN to tape at FNAL. The storage element functionality is expected to be deployed in the LCG-2 release. We expect at next week's CERN demonstrations to show data moving between storage elements.

ATLAS simulates Higgs on Grid3!

Using tools from the Chimera/Pegasus and Globus RLS groups in an integrated Grid Component Environment (GCE) developed by the ATLAS team, 748 simulation and reconstruction jobs were run on fifteen Grid3 sites and delivered to a data cache at Brookhaven National Laboratory. The dataset was cataloged and analyzed using DIAL to produced four electron invariant mass plot

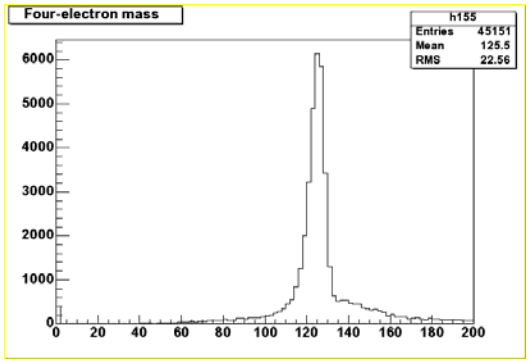

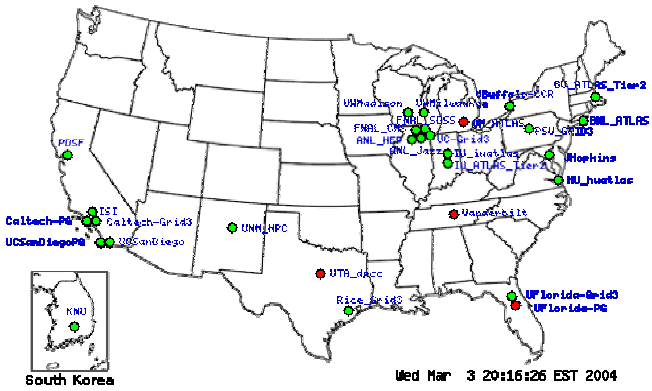

Several documents written to support the development and operation of Grid3 are posted at
http://www.ivdgl.org/grid2003/documents/document_server/show_docs.php?series=grid3&category=tech

Grid3 2004-19 The Grid2003 Production Grid:Principles and Practice

Grid3 2004-16 Grid3 HPDC Abstract

Grid3 2004-15 Grid3+ Site Charter v8

Grid3 2004-14 Grid2003 Installation Howto

Grid3 2004-13 Grid3 monitoring metrics and architecture

Grid3 2004-12 Grid 2003 User Registration And VO Management Policy

Grid3 2004-11 Grid3 Pacman Package: Structure and Implementation

Grid3 2004-10 Grid 2003 Glossary

Grid3 2004-9 Grid2003 Site Policy and Procedures

Grid3 2004-8 Grid2003 monitoring: installation/configuration/troubleshooting

Grid3 2004-7 Site Authentication with the Grid3 VOMS Infrastructure

Grid3 2004-6 Setting Relative Priorities of Testbeds on Batch System

Grid3 2004-5 Simple installation of VOMS

Grid3 2004-4 Grid3 Naming Conventions

Grid3 2004-3 Monitoring HOWTO

Grid3 2004-2 Grid3 Information Directory Schema

Grid3 2004-1 Grid2003 Project Lessons (version 8)

Grid3 2003-3 Grid3/Grid2003 Planning document: Version 21

Grid3 2003-2 GIS (Grid Information System) Configuration

Grid3 2003-1 Tips for Making Good Pacman Caches

A Lessons Learned document was written and reviewed by the Stakeholders. PPDG-37.

An abstract and paper for HPDC-13 were written and submitted. The abstract reads: “The Grid2003 Project has deployed a multi-VO, application-driven grid laboratory capable of sustaining production level services required by the physics experiments of the Large Hadron Collider at CERN, the Sloan Digital Sky Survey project, the gravitational wave searches by the LIGO experiment, the BTeV experiment at Fermilab, as well as applications in molecular and genome analysis. The deployed infrastructure has been operating since November 2003 with 27 sites, a peak of 2800 processors, work loads exceeding 1300 simultaneous jobs, and data transfers among sites of greater than 2 TB/day. We describe application requirements and resulting principles adopted for deployment and configuration of grid services, monitoring infrastructure, resulting application performance and metrics, and operation experience including a summary lessons learned.”

The common grid environment continues to be used by the stakeholders.

3.3 Global Grid Forum and PNPA Research Group

The PNPA-RG started the first deliverable a draft informational document that defines the requirements of the HENP community at this time from GGF WG and moved on with organizing of the workshop to be held at GGF-10. A successful workshop was held with over 65 participants. The agenda and presentations are posted at http://www.ivdgl.org/events/view_agenda.php?id=94. A summary will be posted off the PNPA-RA web pages at https://forge.gridforum.org/projects/pnpa-rg/

3.4 Joint Technical Board (JTB)

The Joint Technical Board held 2 phone meetings. The main content was to proceed with a plan for addressing the divergent and incompatible versions of RLS in the U.S. and Europe. Experiment liaisons were identified as requested for this Joint Project. In the end, however, no effort was available for this and the project has languished.

4 Single Team Reports

4.1 BaBar

Part of the SLAC BaBar group has continued its efforts in putting the SRB into place as the production tool for distributing BaBar data from Tier A to Tier A: SLAC to ccin2p3.

The group successfully demonstrated data distribution of BaBar Objectivity Conditions data from SLAC to: ccin2p3, Iowa, Bristol, Rome during Super Computing 2003. The demonstration used SRB v2.2 and ran continuously for 1 week without major problems. This demonstration provided the SLAC group with

valuable experience in running the SRB system continuously with clients from different countries. This experience will be used in developing tools to provide a robust and reliable system.

Since then the group has worked on the deployment of SRBv3 along with a local set of tools to move production data from SLAC to ccin2p3. Two production MCAT's were established at SLAC. The first one contains the root-files that store the BaBar event data and the second MCAT holds conditions databases and background trigger data.

Work is progressing on the integration of the SRB with the BaBar tools that handle data imports from production centers in Padova and GridKA. These tools populate the SRB MCAT at the same time as the data is backed up to the SLAC HPSS system making the data readily available for export to ccin2p3.

The groups current focus is on the production Tier A to Tier A distribution tool as well as a tool for distribution of conditions data from Tier A to Tier C centers. BaBar root files were registered in the SRB and the transfer to ccin2p3 was successfully tested. We are ready to transfer all run1,2 and 3 data to ccin2p3 which are about 25,000 files with a total size of about 15TB.

JImport is not used for the import of the root event data and therefore the integration with the SRB transfer is not pursued anymore. However, for the Objectivity conditions data the integration of JImport and to copy the data using the SRB has been completed at ccin2p3.

The group gave two presentations at conferences on the work done:

IEEE, Oct19-25 2003, Portland, Oregon
http://www.slac.stanford.edu/BFROOT/www/Computing/Offline/DataDist/DataGrids/srb/talks/ieee2003wk.pdf

ACAT03, Dec01-04 2003, KEK, Tsukuba, Japan
http://www.slac.stanford.edu/BFROOT/www/Computing/Offline/DataDist/DataGrids/srb/talks/acat03ah.ppt

4.2 CMS

Production GRID support and operation remained main focus of CMS-PPDG effort. After successful completion of 50 Million Simulation Events during summer, CMS Pre Challenge Production entered into OSCAR production (Hit Formatting). This new POOL based Hit Formatting tool was quickly made to run over GRID via MOP and McRunjob modifications for MOP.

The USCMS Production GRID became the exclusive choice for producing 16 Million OSCAR Events, with an estimated 5 to 6 months of non-stop. US-MOP GRID production tools went through some feature improvement and maintenance to include direct d-cache storage and retrieval of data files, at worker nodes, and or through cmsgridftp gate. USCMS Production GRID made opportunistic use of GRID-3 resources, and at times over 1100 CPUs were used from two or more submission sites in parallel. No major breakdowns or scalability issues encountered during current operations and events were delivered to CERN at a steady rate. The manual effort to run this large scale production reduced to almost ~1.5 FTE effort. Adding worst-case disk fill ups and other scenarios, the total percentage of successful jobs around 70%, with best organizations and planning a 90 to 100% successful turn-around was achieved.

CMS-PPDG MOP continued to look at making Distributed Production Environment tools more operators friendly and reduce debugging and troubleshooting effort, even further. Condor_G match making is made integral part of GRID submission to avoid scheduling mechanism, though manual scheduling is still a proffered method. Evolving DPE tools into GRID/Web-Services, integrated via CLARENS was successfully prototyped. Work is under prototyping phase to create a logging and tracking database based GRID submission and troubleshooting tool with simple to use Graphical User Interface. This might, will also facilitate to run private production or user analysis over GRID. Digitization has recently been started at CMS, and will be run over Production GRID very soon; necessary modifications in GRID production tools (MOP/McRunjob) are already tested and ready to go into latest version of DPE (2_4_1). New POOL based Digitization, catalog operations made it little difficult to run Digitization over GRID.

The CMS Data Challenge 2004 has recently started. DC04 is giving a better understanding of Batch and User Analysis model for CMS. Once Local Farm based Analysis tools are satisfactory ready for operations,

GRIDification will be achieved. The architectural changes introduced in McRunjob due to requirements posed by GRID/MOP, have made it much easier and smoother to adopt Production/Analysis jobs for GRID.

The CMS Caltech group continued to work in the area of high performance networking and distributed analysis and monitoring. They attended World Telecom 2003 exhibition in Geneva, where a Grid-enabled analysis demonstration using Clarens was given as part of the program at the CERN booth. Two major demonstrations of GAE was given at the Supercomputing 2003 conference in Phoenix, AZ. An invited tutorial on the Clarens Web Services framework was given at the LISHEP 2004 Digital Divide workshop at the State University of Rio de Janeiro (UERJ) in February 2004. A demonstration of high-speed disk-to-disk data transfer using Clarens, with a peak of 400 MB/s (3.2 Gb/s) bandwidth over a 10 Gb/s connection using two high-capacity disk servers at the CENIC network POP in Los Angeles, coupled with two similarly configured Clarens client machines on the show floor in Phoenix. Also a distributed analysis demonstration in collaboration with the University of Florida was given at the Caltech CACR booth using Clarens as a communications layer for different components of the demonstration.

Clarens: Two new releases were made during this period, the latter the first to support access control for files, as described during the preceding reporting period. A single-step installation script was developed for sites that require only client functionality. A server and client distribution for use in the ATLAS experiment was developed in collaboration with Miguel Branco from that experiment. The complete method documentation on the Clarens server was updated and extended with descriptions of new functionality added, described below. A printed manual was also started that includes installation, configuration, service writing, client usage as well as authentication protocol information. Implemented a service lookup and registration service that allows authenticated third party registration of services. The Java Clarens server was further developed with functionality close to that of the Python/C version. Development of a database-backed logging facility for web service requests and responses was started that assists in tracking data provenance. Logging can occur at three levels of detail, 0 for no logging, 1, for method name, time and user logging, and 2 which also stores request parameters and return values. Development of a configuration monitoring service was started, which stores key/value pairs organized in user-defined categories. This is an oft-requested service that is used to store and retrieve configuration parameters of Grid sites, e.g. paths to libraries and executables. Started a Dataset Catalog Service based on the API from the PPDG CS-11 group.

Monitoring: MonaLisa is used as the main system monitoring and presentation tools in Grid2003: The MonaLisa database was extended and used by the University of Chicago grid telemetry project to provide ongoing metrics for Grid3.

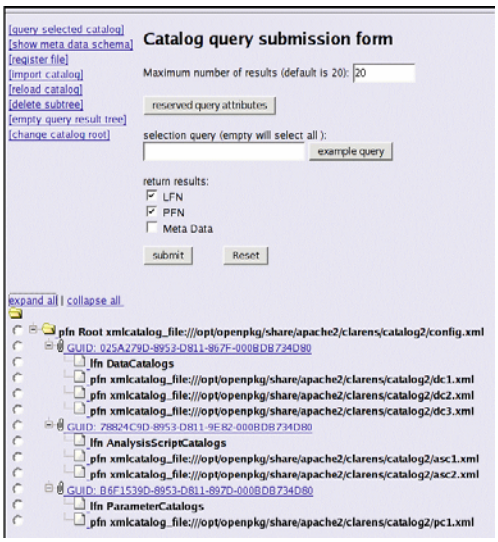

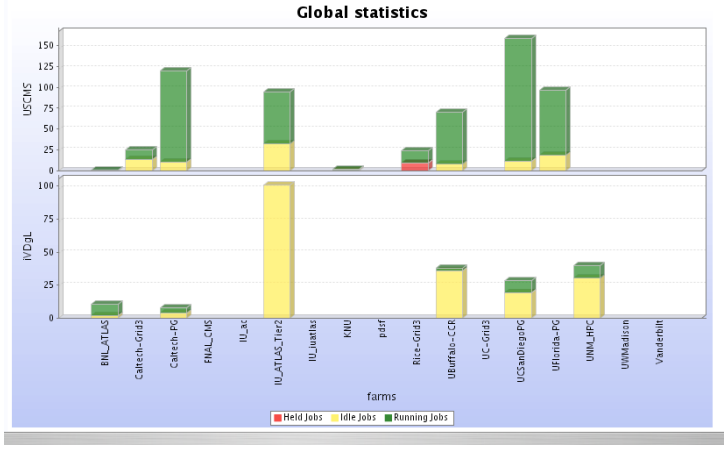

4.3 D0

D0 continued to deploy, harden and robustify JIM job management component into the main experiment monte carlo simulation production systems. This was slow and very detailed work with each of the sites participating in the deployment. Much progress was made and JIM is now regarded as a legitimate service which is expected to be more widely used over the next few months.

The SAM data management and grid system continued to be used for production across 20 some sites in D0 for the analysis and reprocessing of the experiments data. Metrics for many different performance characteristics are being kept.

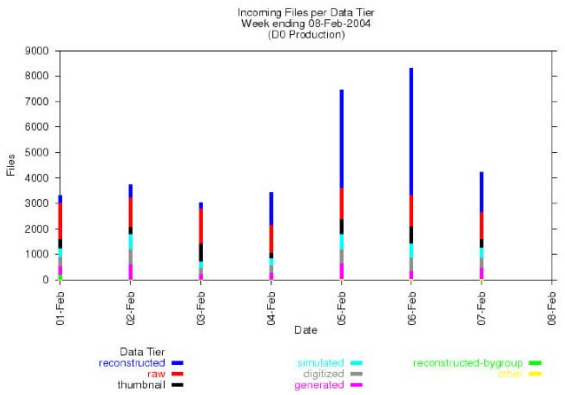

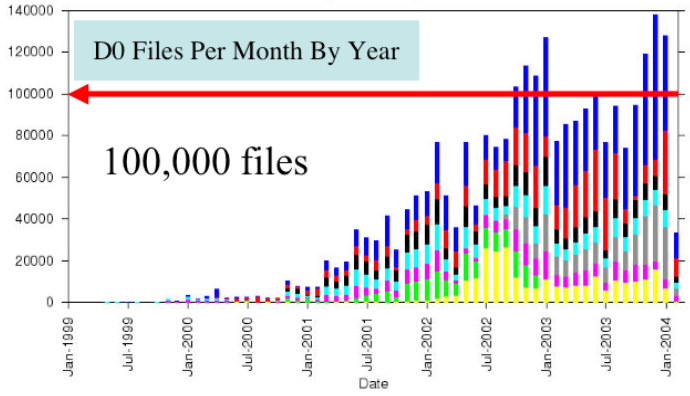

Success/failure metrics are an important component of the JIM deployment. The goal is to have application use of the Grid approach the success rate of the application run on a local cluster under SAM management, which is about 94%. Both the JIM infrastructure and the end to end efficiency of running the application over the grid on three sites were much improved.

	Manchester	CCIN2P3	Wisconsin

Jim infrastructure efficiency	>99%	>99%	>99%
End to end efficiency of running the application – includes Site efficiency x Code efficiency	~85%	~60%	~60%

4.4 JLab

During this quarter we suffered setbacks resulting from Hurricane Isabel. Our focus was directed toward the immediate needs of the lab and getting the facility up and operational.

We continued to participate in discussions surrounding the final revision on the SRM v2.X specification. Most discussions were focused on interpreting the written specifications for implementation, particularly in the area of SURL interpretations for locating files. The SRM specification itself does not contain the WSDL for the web services. However, the SRM collaboration is developing the WSDL for use in SRM development. A preliminary WSDL was evaluated and our comments were returned for consideration. In February a version of the WSDL was generated, and we are now in the process of adapting our implementation to this specification, as are other members of this collaboration. Minor implementation issues are being addressed as we progress, with standardized implementations expected in the coming quarter. Jefferson Lab has begun implementation of phase 1 for a prototype of a Grid Job Scheduler using web services. Four phases of development are currently planned for the Grid Job Scheduler. Phase 1 consists of a client translator, a job scheduler with a defined interface and WSDL for job submission, a site submitter, and job monitoring capabilities. The client translator interprets a job description from a user and submits it to the job scheduler. The job scheduler decides which site the job or jobs should be dispatched to for execution. The site submitter interfaces between the job scheduler and the batch system (LSF, PBS, Condor) for a site. This work is in the context of a collaboration with STAR, and like the SRM work will support both the experimental physics program and the Lattice QCD program at Jefferson Lab. Due to the hurricane, our schedule for phase 1 deployment has slipped by 4-6 weeks.

As part of our work on the Grid Job Scheduler, we are collaborating with STAR on a user job description language (uJDL or RDL for Request Description Language). We are working for a solution, which could be common across multiple experiments. It is envisioned that both STAR and the JLab Grid Job Scheduler will make use of this common uJDL. Current planning has the integration of the uJDL with the Jefferson Lab Grid Job Scheduler during phase 3 of its development.

4.5 STAR

Gabriele Carcassi has left STAR for a new position while Valeri Fine and Lidia Didenko have joined the effort as per the development, testing and consolidation of the Grid infrastructure and core-STAR software. Due to the PPDG team reshape and the temporary technology transfer overhead consequent to Gabriele's leave, we expect our deliverables to shift toward a more STAR centric consolidation (were the current and immediate expertise is) and a set back in deliverables oriented toward providing community work. We would like to thank Gabriele for his efforts and valuable contributions and wish him success in his new endeavors.

4.5.1 Network

Our data transfer from BNL to NERSC close to 30 MB/sec from over 50 streams. This data transfer rate is, to date, the best reported, thanks to improvements made by Dantong Yu (RCF) through upgrades of the network interface. However, concerns arose as per the drastic difference between the single stream and multi stream performance, a single stream pushing around 1 MB/sec through the same path. This drop of performance was believed to be due to a drop of packets at the BNL firewall. With the assistance and expertise of the BNL Information and Technology division (ITD) networking group, several network tests were conducted as per identifying possible bottleneck on the BNL network. Those tests were made to measure the performance for a transfer from inside the multiple firewalls to the perimeter firewall (just

before the data is shipped out). A rate of 430 Mbits/sec was measured which seems to indicate that the possible loss in performance throughout the path to NERSC is not due to an internal BNL network problem (or drop of packets between inside and the BNL perimeter). This investigation will have to be resumed by attempting similar tests on the NERSC side.

We would like to thank the ITD networking group of Scott Bradley for their effort and assistance in this matter. Special thanks go to John Bigrow and Terry Healy for their time and patience. Equal thanks to Dantong who has provided support for this activity.

4.5.2 File transfer and data management

Production quality grid based file transfer continued stable running throughout the period. The data rate and transfer are represented on the graph above.

During this period, we continued the use of HRM for file transfer between BNL and NERSC/PDSF. Additionally, a new component of the Replica Registration Service (or RRS) was tested in production mode (with recovery for minimal impact on transfer). Over 25k files were transferred using RRS, amongst which 500 missed the registration due to miscellaneous problems we will resolve and report later. Considering the small percentage of the current failure (i.e. 2%), we consider this exercise a success and are very encouraged by he approach. The Replica Registration service was implemented as a user-plugin calling a user registration script interfacing with the STAR FileCatalog. To date, 4 millions replicas are registered in our catalog representing 3 million files. The current year-4 run is foreseen to increase our catalog to 5.5 million primary files with a minimum of 7 million replicas.

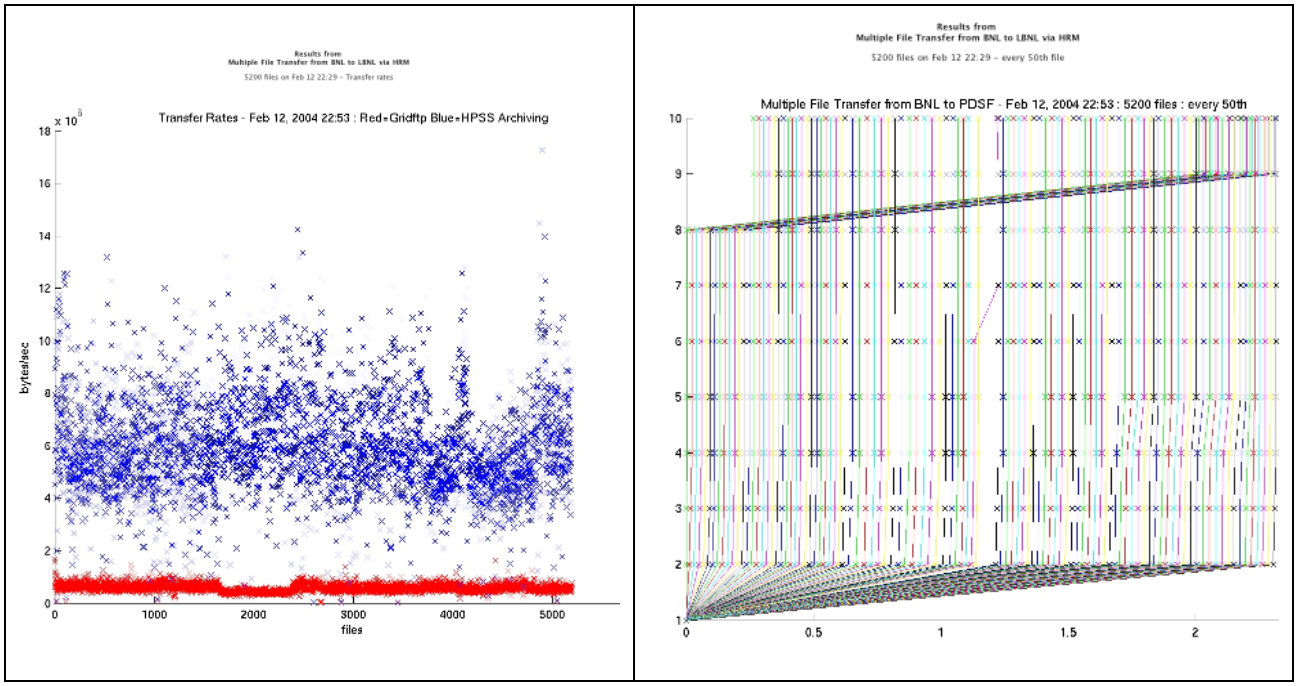

4.5.3 Monitoring

This effort is supported by Efstratios (Stratos) from the BNL/ITD. Stratos has been working on developing a MonaLisa module aimed at queue monitoring. Based on the **GLUE Schema** and especially on the CE information provider, a MonaLisa module was written to provide a consistent view of queue activities. The scope of this work was presented at the last PPDG collaboration meeting. Currently implemented and tested for Condor, PBS and LSF, we deployed this module at our two main sites. We have however

contacted the MonaLisa developers has recent upgrade in the MonaLisa framework has introduced a glitch in this module. We hope to incorporate our current work as yet another community usable component.

A MonaLisa module was also developed to provide to the STAR Scheduler the information collected by the load information (Ganglia). This will later be extended to incorporate the queue information and provide the basis for a more complex scheduling policy.

4.5.4 Grid job submission

During this period, tremendous progress has been made on several fronts. First, earlier in this period, Gabriele Carcassi provided a patch to the LSF.pm module (see **Bugzilla #950**, follow-up **83**, **84**, **85**) resolving a problem we were experiencing while using LSF as the job manager back-end. The symptoms were that while LSF would consider some jobs pending, globus would report those same jobs as pending. After investigation, it appeared that the logic in LSF.pm was such that a job would be submitted and LSF immediately polled for information. The delay between the submission time and the report from polling the queue appears to be specific to a multithreaded LSF environment. A patch was subsequently delivered against the VDT 1.1.12 (Globus 2.4.2). Note that this effort alone brought our job submission success rate of 100% to 99.9% (one of the test series showed a single job unrelated failure) and all successful jobs now return their output. A consolidation of this patch will be required in future.

The STAR scheduler was extended to implement and introduce more flexibility as per running simulation jobs. In this mode, one job description is expanded to many similar jobs; our previous implementation considering splitting of a large dataset (and its availability) as the primary condition for job splitting. Also, the current schema for the U-JDL has been extended to incorporate action to be performed on the resulting outputs. For example, after an output is created, registration may occur, a copy left on a mass storage, a globus-url-copy executed or all of the above. The **current schema** is represented on our documentation pages as per the pre-common-project RDL definition.

Further enhancements interfaced the Scheduler to a MonaLisa module providing information on the load for diverse clusters. Our first trial is to test its feasibility at local level, that is, use two separate sub-clusters accessible via two different queue managers and understand better if an automated load balancing can be achieved with this simplistic approach. The clusters used currently are our user analysis farm and/or (usually excluded from user access) reconstruction farm. This was integrated within a separate job submission policy to avoid interfering with on-going work and usage. Our phase 2 will incorporate the queue information and define a better submission policy; phase 3 will take advantage of our discoveries for Grid job submission (dynamic throttle to a given site) shall the approach be a success. Only local job submission to independent clusters has been tested, seems as providing an improvement as per the discovery of idle resources, but we do not have enough statistics (or run long enough) to report definite success and/or benefits.

Part of the recent developments also released a Grid job submission policy which was tested to simulate 1000 jobs of 1000 Hijing events each. Initially hectic and unstable, we found Grid job submission depends on lots of external factors (including NFS stability, load on the gatekeeper, etc). In optimal conditions however, we measured a 98.7% success, the missing jobs were aborted for unknown reasons. Those tests were conducted primarily by Lidia and Valeri. As for the LSF module, we will investigate this in depth, report to the community and possibly provide solutions were applicable. As a side note, while encouraging, there are however numerous remaining problems we have not yet addressed (but aware off) such as the submission of a large number of jobs on the grid and the port range availability. It is our belief and our feedback as per our needs for job management performance, that a system must be able to handle up to 100k jobs a day sustained to be usable. In fact, the STAR user community dispatches approximately 30-50 k jobs a day and this amount is seen to rise rapidly with the venue of a factor of x5 in data amount.

Finally, and on another front, we have continued work with Jefferson Lab team in the definition of U-JDL and RDL foreseen as a common ground across experiments to describe jobs or tasks to be executed on the grid as well as a certain level of tracking and control. Our common project was well stated and a first version of the RDL was published as **PPDG-39**. We received feedback and suggestions through the PPDG job management list but have also received positive feedback from a few teams. We will resume this effort

with JLab and expand the current documents available on our Request Definition Language documentation **page** as soon as possible (timing was delayed by both team experiencing temporary set-backs).

4.5.5 Education, outreach and collaborative projects

We continued support for the BNL Technology Meeting. As examples, Ruth Pordes presented an overview of the "**Grid2003 – A Multi-Organization Functioning Grid Demonstrator**" and Greg Nawrocki presented a "**The Globus Toolkit and the OGSI – WSRF Evolution**". Other talks covered topics such as Web Services and Highlights of the Globus World 2004. The detailed content may be seen from our **BNL Technology Page**. We would like to thanks all the speakers who have agreed to make presentation throughout the years. At this stage of maturity of this meeting, speakers being available from remote (thanks to Tameka Carter and ITD for helping with the introduction and handling of the AccessGrid system), we have asked for a new candidate to take on the organization logistic. Dantong Yu from the RCF has agreed in pursuing this effort.

During this period, continuous effort to converge on a similar approach has been made with the Phenix collaboration. Joint meetings are regularly scheduled: the STAR Scheduler has been of interest to Phenix and consequently modified to suit the need of both collaborations. Our knowledge in monitoring (MonaLisa and beyond) has also been greatly beneficial to Phenix and in this area, we have provided technical assistance and shared our vision as much as possible. We would have hoped to achieve a similar outreach to the Atlas collaboration but have not succeeded to date to define common grounds and/or projects. A soon coming Grid workshop at BNL through the iVDGL collaboratory project will be a last opportunity to re-iterate this effort.

4.6 Condor

The Wisconsin group provided many experiments and projects (including Grid3, USCMS, DZero/SAMGrid, CDF, and STAR) with design, development, integration and support services for VDT, Condor, DAGMan, Grid Shell, Grid Exerciser, etc. Peter Couvares served as the Job Management Working Group co-chair.

The Wisconsin team also contributed to the Data Management area, investigating integration of NeST and SRM, as well as supporting and packaging of the VDT.

The Wisconsin team provided effort for troubleshooting, the Grid Exerciser application and application deployment on Grid2003.

4.7 Globus - ANL

4.7.1 Coordination and Support

Continuing interactions in terms of coordination and support of the PPDG applications included weekly phone meetings and email lists for Atlas and CMS, following the grid emails lists of D0, and providing support for the Argonne-Chicago ATLAS team in their efforts to perform "data challenge on demand" event generation using VDT, RSL, and Chimera. Support through the discuss lists and Bugzilla are available to all experiments.

4.7.2 Globus Toolkit 2.x updates and bug fixes

We resolved 3 bugs listed in Bugzilla (260, 1283, 1284) and have no known open PPDG-related bugs still in our system. A patch has recently been posted for Bug 950, however, this was not considered to be a bug by the Globus Toolkit team but rather an anomaly that occurs as the result of an inconsistency within certain implementations of NFS (full report available upon request). Additional information about Bugzilla bugs can be found at http://bugzilla.globus.org

4.7.3 Globus Toolkit 3.2

The Globus Toolkit 3.2 alpha was made available for download. Although stability improvements were made to the open source implementation of OGSI (several OGSI - compliant services) it is nevertheless an alpha release and should not be used in production settings. The release is available from http://www-unix.globus.org/toolkit/3.2alpha/download.html with additional advisories posted in the release notes: http://www-unix.globus.org/toolkit/3.2alpha/release-notes.html

The Globus Toolkit 3.2 beta is available for download. This release contains important bug fixes and improvements on existing functionality. Changes include better deactivate semantics, a change in the persistent grid service programming model, changes to the security descriptor format, updated third party jars and more. The release is available from http://www-unix.globus.org/toolkit/3.2beta/download.html with additional advisories posted in the release notes: http://www-unix.globus.org/toolkit/3.2beta/release-notes.html

4.7.4 GridFTP / RFT

The wuftpd based GridFTP server and globus-url-copy were released with the GT3.2 alpha and beta releases. They contain the feature set listed in the Q3 report. XIO is also present in 3.2. Additional information on XIO can be found on a newly added section to the Globus Toolkit web site at: http://www-unix.globus.org/developer/xio/

A bare bones version of the new (non-wuftpd based) GridFTP server was released with the 3.2 alpha, but will be removed from the 3.2 distribution due to the accelerated release schedule brought on by the WSRF announcement. We did use the new server code base at SC2003 in the bandwidth challenge and achieved nearly 9 Gbs with 15 servers running over the TeraGrid network. The new server will be feature complete by the end of March 2004 (non-striped) and will be available for limited distribution testing.

RFT has also had significant improvements made in terms of functionality and scalability. RFT can now accept a directory and it will move the entire directory tree. Scalability has been improved so that RFT can now handle approx 4000 files. We are working on further scalability improvements. Our goal is to be able to handle millions of files.

4.7.5 Monitoring and MDS work

The MDS2 scalability analysis work continued, concentrating on adding netlogger calls to the MDS infrastructure to determine bottlenecks. This work was submitted and accepted for publication ("Performance Analysis of the Globus Toolkit Monitoring and Discovery Service, MDS2", Xuehai Zhang and Jennifer M. Schopf, to appear in International Workshop on Middleware Performance (MP 2004) in connection with the 23rd International Performance Computing and Communications Conference (IPCCC), April 2004.)

While current monitoring plans involve MDS2, the evolution of the PPDG will require the capabilities of the web services-based MDS3 (OGSI-compatible and part of the Globus Toolkit version 3 release) and potentially MDS4 (WSRF-compatible and part of the upcoming Globus Toolkit version 4 release), for its full deployment. MDS3 and MDS4 both represent a re-architecting of the MDS. While MDS3 does not, as yet, present radically new functionality to the end user, it lays the foundation for new features, such as registration, which will allow entities to request notification of changes in information instead of having to poll for such changes. The Globus Toolkit team will continue work on an MDS3-compatible archiver, which should be available in the June 2004 time frame, and which will then be made WSRF-compliant.

4.7.6 Security

The 3.2 alpha release includes a number of components brought about by ANL's work with PPDG. This release includes the authorization callout for the 2.4-compatible portions of GT. This callout is already in use at FNAL to enable the integration of their Site Authorization System (SAZ) with their Globus deployment.

Work in the GGF OGSA-AuthZ working group continues to standardize an equivalent authorization callout for use in WSRF-based portion of GT4. This standardization effort is greatly helped by the experience gained during the Site-AAA work and design of current authorization callout .A production release of CAS, implemented as an OGSI service, is included with the GT3.2 releases. Additional information on CAS can be found at http://www.globus.org/Security/CAS/GT3/

4.7.7 Training, Presentations and Papers

The following Globus related presentations and tutorials were given at the 2003 Supercomputing Conference. See the Supercomputing Conference web site http://www.sc-conference.org/sc2003/ for more details.

1. The Grid: Software Standards for Cyberinfrastructure, November 16, 8:30AM -12:00PM, Carl Kesselman (USC Information Sciences Institute)

2. How to Build a Grid Service Using the Globus Toolkit ® 3 November 16, 8:30AM -5:00PM or November 17, 8:30AM 5:00PM (the latter is already sold out), Lisa C. Childers (Argonne National Laboratory), Charles A. Bacon (Argonne National Laboratory), Ravi K. Madduri (Argonne National Laboratory), Ben Z. Clifford (The Globus Project)

3. A Tutorial Introduction to High Performance Data Transport, November 16, 8:30AM 5:00PM, Bill Allcock (Argonne National Laboratory), Robert Grossman (University of Illinois at Chicago), Steven Wallace (Indiana University)

4. Grid Services for Data Management and Virtual Data, November 16, 1:30PM 5:00PM, Ann Chervenak (University of Southern California), Ewa Deelman (University of Southern California), Mike Wilde (Argonne National Laboratory)

GlobusWORLD 2004 took place on January 20 – 23 in San Francisco, CA. The Globus Alliance and IBM in conjunction with HP announced details of the new WS-Resource Framework standard, a further convergence of Grid services and Web services. Additional information regarding the conference can be found on the GlobusWORLD web site at: http://www.globusworld.org. Conference presentations and abstracts, including those specific to the WSRF announcement, can be found via the hyperlink conference program at: http://www.globusworld.org/program/conference.asp. Additional information regarding the workshops, including select abstracts and presentations can be found at: http://www.globusworld.org/program/workshops.aspSRM

HRM was deployed at PDSF and BNL with new features including directory browsing and hierarchical directory-directory replication (joint work with Junmin Gu). A new file monitoring tool that works from behind the firewall (joint work with Viji Natarajan) was developed and deployed. A demonstration of all these capabilities was shown at the SC 2003 demo at Pheonix AZ

The first version of RRS was completed and deployed on PDSF (development done by Viji Natarajan) and will be tested by STAR.

The DRM-WSG and HRM-WSG software is being readied for VDT packaging as web services gateways in joint work with Junmin Gu and Viji Natarajan.

4.8 SRB

The SDSC team released version 3.0.1 of the Storage Resource Broker in December. This includes Peer-to-Peer federation of MCAT servers, a feature being used by BaBar. A web interface and Java interface have been ported to SRB version 3. The Java administration tool has been enhanced for simplified administration of zones (independent data grids). This work is described in a presentation at the December collaboration meeting, http://www.ppdg.net/mtgs/15dec03-lbnl/ppdg-srb3-matrix.ppt.

Multiple reports were written on distributed data management for the Data Area of the Global Grid Forum, and other conferences:

1. Jagatheesan, A., R., Moore, “Data Grid Management Systems,” NASA / IEEE MSST2004, Twelfth NASA Goddard / Twenty-First IEEE Conference on Mass Storage Systems and Technologies, April 2004.
2. Moore, R., “Preservation Environments,” NASA / IEEE MSST2004, Twelfth NASA Goddard / Twenty-First IEEE Conference on Mass Storage Systems and Technologies, April 2004.
3. Malaika, S., A. Merzky, R. Moore, "GGF Data Area Structure and Function Analysis", Global Grid Forum 10, Berlin, March 2004.
4. Moore, R., “Evolution of Data Grid Concepts”, submitted to the Global Grid Forum Data Area Workshop, January, 2004.
5. Rajasekar, A., M. Wan, R. Moore, W. Schroeder, “Data Grid Federation”, submitted to PDPTA 2004 - Special Session on New Trends in Distributed Data Access, January, 2004.
6. Moore, R., “Operations for Access, Management, and Transport at Remote Sites,” submitted to the Global Grid Forum, December, 2003.
7. Moore, R., A. Merzky,“Persistent Archive Concepts,” submitted to the Global Grid Forum, December, 2003.

4.9 IEPM, Network Performance Monitoring

4.9.1 Web/Grid Services

The new IEPM-BW Oracle backend database now provides easier ways to search through the IEPM-BW data. We followed this up by providing Grid services access to the data utilizing the GGF NMWG naming recommendations. To demonstrate the utilization we created a web page to interactively access the IEPM-BW data via the Grid services interface. More information on the Grid services access can be found at the IEPM Web Services page. Our next steps will be to provide similar access to the PingER data. Lee and Paul Meallor) to discuss and understand the OGSI schema. Following this Warren worked on porting the MAGGIE Web service to an OGSI type grid service using the perl module. We extended our traceroute visualization tool to add: drill down to view available bandwidth time-series for the selected path; and to allow navigation through the historical archives of traceroutes. We wrote, submitted and had accepted an extended abstract short abstract on *Correlating Internet Performance Changes and Route Changes to Assist in Trouble-shooting from an End-user Perspective*, that describes this work.

4.9.2 PingER

There is increasing interest from the HENP and other scientific communities to understand and do something about the Digital Divide, i.e. the difference in Internet performance to developing and developed countries. Since 10-20% of HEP collaborators on the major experiments come from countries in developing nations, this is very important to HENP. We prepared the *January 2004 Report of the ICFA-SCIC Monitoring Working Group* for ICFA. In preparation for the WSIS in Geneva and the ICFA/SCIC Digital Divide meeting in Rio de Janeiro in Feb 2004, we extended the number of remote sites monitored by about 10%. This was focused on providing better information on the Internet performance to developing countries. We also worked with NIIT Pakistan to successfully install a monitoring node there. This brings the number of countries with monitoring hosts to 13. Early results for NIIT were provided to the director of NIIT for a presentation recommending upgrading their connectivity.

4.9.3 High speed networking

We attended and exhibited our monitoring projects at the SLAC/FNAL booth at SC2003. In addition, we once again participated in the Bandwidth Challenge. SLAC's challenge, in collaboration with Caltech and LANL, this year was Distributed Particle Physics Analysis using Ultra High Speed TCP on the Grid and it captured the SuperComputing 2003 Bandwidth Challenge award for the most data transferred (6.6TBytes in under 50 minutes).We co-authored a paper with LANL and Caltech on Optimizing 10-Gigabit Ethernet in Networks of Workstations, Clusters, and Grids: A Case Study that was accepted and presented at

SC2003. As co-chair of the 2nd Protocols for Fast Long Distance Networks to be held in Chicago in February 2004, we organized a review team, and reviewed and selected papers to be presented. We started to work on automated fault finding by looking at the IEPM-BW monitoring data. We presented *Measurement and fault finding using MAGGIE and PIPES* at the Internet2 Member meeting, Indianapolis Oct.13-16.

5 Appendix

5.1 List of participants

TEAM	Name	F	Current Role CS	Systems and Production Grids	Job Mgmt	Data Mgmt	AAA	Grid Analysis and Catalogs	Other: Web Services, Evaluations Interoperation, etc.
Globus/ANL	Ian Foster	Y	Globus Team Lead, GriPhyN PI, iVDGL, GriPhyN			x			
	Mike Wilde	N	GriPhyN coordinator, ATLAS-CS liasion	x		x			
	Jenny Schopf	N	GriPhyN, iVDGL, liaison,	x		x			x
	Greg Nawrocki	Y	PPDG Globus liaison	x		x			x
	Willliam Allcock	Y	GridFTP	x		x			x
	Von Welch	N	CAS				x		
	Stu Martin	N		x	x				
ATLAS	John Huth	N	ATLAS Team lead, GriPhyN Collaborator	x					
	Torre Wenaus	N	LCG Applications liason		x	x		x	
	L. Price	N	Liaison to HICB, HICB Chair						
	D. Malon	N	Database/POOL Liason					x	
	A. Vaniachine	N						x	
	E. May	N	Testbed applications	x		x			
	Rich Baker	N	Testbed applications, VO tools	x			x		
	Kaushik De	N	Testbed applications	x					
	David Adams	Y	Distributed analysis					x	
	Wensheng Deng	Y	Metadata catalogs			x		x	
	R. Gardner	N	IVDGL coordinator, Atlas Grid Tools		x	x			x
	G. Gieraltowski	Y	Interoperability	x				x	x
	Dantong Yu	Y	Monitoring and VO	x			x		
BaBar	Richard Mount	N	PPDG PI, BaBar Team co-Lead						
	Tim Adye	N	BaBar Team Co-Lead						
	Robert Cowles	N					x		
	Andrew Hanushevsky	Y				x			
	Adil Hassan	Y				x			
	Les Cottrell	N	IEPM Liaison	x					
	Wilko Kroeger	Y				x			
CMS	Lothar Bauerdick	N	CMS Team Lead. GriPhyN collaborator						
	Harvey Newman	N	PPDG PI. GriPhyN collaborator, Co-PI iVDGL						
	Julian Bunn	N	CMS Tier 2 manager, GriPhyN & iVDGL collaborator	x				x	
	Conrad Steenberg	Y	CS-8:Analysis Tools, GriPhyN collaborator					x	x
	Frank Lingren	Y	CS-8:Analysis Tools, GriPhyN collaborator					x	x

	Iosif Legrand	N	CS-8:Monitoring Tools						x
	Vladimir Litvin	N	GriPhyN collaborator		x				
	James Branson	N	CMS Tier 2 manager	x					
	Ian Fisk	N	CMS Level 2 CAS manager, iVDGL liaison	x					
	James Letts	Y	Working on VDT testing scripts	x					
	Saima Iqbal	N	web technology evaluation					x	
	Suresh Man Singh	N	grid deployment	x					
	Anzar Afaq	Y		x	x			x	
	Greg Graham	N		x	x			x	
Coordination	Ruth Pordes	Y	PPDG coordinator			x			
	Doug Olson	Y	PPDG coordinator			x	x	x	
	Miron Livny	Y	PPDG coordinator		x	x	x	x	
	Joseph Perl	Y	CS-11 co-coordinator					x	
D0	Wyatt Merritt	N	D0 team Lead	x	x				
	Igor Terekhov	Y	JIM Team Lead	x	x				
	Andrew Baranovski	Y		x					
	Gabriele Garzoglio	Y		x	x	x			
	Sankalp Jain	Y	Through contract with UTA CSE Department	x	x				
	Aditya Nishandar	Y	Through contract with UTA CSE Department	x	x				
HRM/LBNL	Arie Shoshani	y	SRM Team Lead. GriPhyN collaborator			x			
	Alex Sim	Y				x			
	JunminGu	Y				x			
	Viji Natarajan	Y				x			
SRB/UCSD	Reagan Moore	Y	SRB Team Lead. GriPhyN collaborator			x			x
	Wayne Schroeder	Y	CS-8: Web Services			x			x
JLAB	William Watson	Y	JLAB Team Lead			x			x
	Sandy Philpott	N	facilities	x			x		
	Andy Kowalski	N				x			
	Bryan Hess	Y	Web Services			x			x
	Ying Chen	Y	Web Services	x		x			x
	Walt Akers	N	Web Services	x		x			
STAR	Jerome Lauret	N	STAR Team Lead	x	x				x
	Gabrielle Carcassi	Y		x	x				
	Dave Stampf	N		x					
	Richard Casela	N		x					
	Eftratios Efstathiadis	N		x					
	Eric Hjort	Y		x		x			
	Doug Olson	N		x		x			x
Condor/U.Wisconsin	Miron Livny	Y	PPDG PI, PPDG Coordinator. GriPhyN collaborator	x	x	x			x
	Peter Couvares	Y			X		x		
	Rajesh Rajamani	N			x				x

	Alan DeSmet	Y			x		x		
	Alain Roy	N			x				
	Todd Tannenbaum	Y			x				
Globus/ISI	Carl Kesselman	N	Globus/ISI lead						
	Ann Chervenak	N				x			
PHENIX	David Morrison								
CDF									
ALICE									

5.2 Year 4 & 5 Milestones (from renewal proposal)

Milestone Summary for Year 4		
May-04	PPDG-CMS	Enable clients to access the Pool metadata catalog via the Clarens web service environment
Jun-04	PPDG-SAMGrid	Test JIM job scheduling at CDF
Jun-04	PPDG-SAMGrid	JIM job scheduling on remote resources for next round of D0 reprocessing
Jun-04	PPDG-SRM	SRM Specification V3
Jun-04	PPDG-BaBAR	SRB based production system for data distribution
Jun-04	PPDG-CS-11	Implementation of the CS11 Dataset Query Service by at least Caltech and SLAC
Jun-04	PPDG-CS-11	Complete specification of 3 more of the standard interfaces for Analysis Services
Jul-04	PPDG-OSG-Authorization	Specification and Requirements for Authorization Version 1
Jul-04	PPDG-OSG-Monitoring	Specify Requirements V1
Jul-04	PPDG-STAR	Grid-based Monte-Carlo production start to migrate from centralized Tier-0 center to a Tier-1 distributed model.
Jul-04	PPDG-STAR	Batch oriented user analysis will start.
Sep-04	PPDG-CMS	Deployment of second version of lookup service for dynamic discovery. Multiple instances of service will be deployed.
Sep-04	PPDG-OSG-Troubleshooting	Specify Baseline Capabilities
Oct-04	PPDG-OSG-Accounting	Specification and Requirements for Accounting Version 1
Oct-04	PPDG-OSG-WMGT	Process Wrapper (ProRap) for remote execution of HENP applications V1
Nov-04	PPDG-SRM	SRM Implementation V2
Dec-04	PPDG-OSG-Accounting	Implementation of Accounting V1.0
Dec-04	PPDG-OSG-Monitoring	Monitoring of the system infrastructure across all the Laboratory facilities
Dec-04	PPDG-ATLAS	Deploy prototype sufficient to assess POOL-based ATLAS applications
Dec-04	PPDG-JLAB	A prototype implementation of SRM, based on specification 2.x,
Dec-04	PPDG-OSG-RRS	Replica Registration Service Specification V1
Dec-04	PPDG-OSG-WMGT	Enhance the scalability of Condor-G (support 105 jobs)
Dec-04	PPDG-SAMGrid	JIM job scheduling as part of the D0 RAC infrastructure
Dec-04	PPDG-SAMGrid	Prototype services which use more general Grid resources, in particular run tests for CDF and D0 on OSG resources
Dec-04	PPDG-CS-11	Multiple Implementations of the three interfaces selected
Dec-04	PPDG-DMGT	Use of Enhanced GridFTP
Dec-04	PPDG-OSG-Monitoring	Instrumentation of selected data transport and replication applications
Jan-05	PPDG-OSG-Authorization	Implementation V1.0
Jan-05	PPDG-OSG	Multi-experiment use of 2000 CPUs, 100TBytes
Jan-05	PPDG-BaBAR	Applications that publish in the RLS the SRB MCAT knowledge on the data location.
Apr-05	PPDG-OSG-WMGT	Process Wrapper (ProRap) for remote execution of HENP applications V2
May-05	PPDG-OSG-Monitoring	Monitoring and instrumentation roadmap
May-05	PPDG-CS-11	Complete specification of remaining ~15 Interfaces

Milestone Summary for Year 5		
Jun-05	PPDG-ATLAS	Access for Grid-based applications to ATLAS metadata services
Jun-05	PPDG-STAR	STAR, data production will be moved to the Grid and a first prototype of an interactive user analysis framework using Grid components
Jun-05	PPDG-SAMGrid	Use remote resources for 50% of CDF user analysis cycles
Jun-05	PPDG-SAMGrid	Begin to add accounting and authorization capabilities to SAMGrid
Jun-05	PPDG-OSG-RRS	Replica Registration Service Implementation V1
Jun-05	PPDG-BaBAR	RLS to make SRB calls to move data from one site to another
Jun-05	PPDG-SAMGrid	SAM operational Model
Jun-05	PPDG-SRM	Specification V3
Jul-05	PPDG-OSG-Troubleshooting	Implementation of Baseline Capabilities
Sep-05	PPDG-OSG-Authorization	Specification and Requirements V2
Sep-05	PPDG-OSG-Monitoring	Updated Specification and Requirements V2
Oct-05	PPDG-OSG-Accounting	Specification and Requirements for Accounting V2
Nov-05	PPDG-OSG-Troubleshooting	Specification of enhanced capabilities
Dec-05	PPDG-OSG-Monitoring	Instrumentation of selected end-to-end data analysis applications consistent with the data transport instrumentation infrastructure
Dec-05	PPDG-DMGT	Use of Enhanced GridFTP
Dec-05	PPDG-OSG-Accounting	Implementation of Accounting V1.1
Dec-05	PPDG-OSG-Monitoring	Enable MonaLisa clients to obtain feedback (through Clarens or other web service infrastructures)
Dec-05	PPDG-SAMGrid	Monitoring and Troubleshooting V1
Dec-05	PPDG-CS-11	First reference implementation of all interfaces
Mar-06	PPDG-SAMGrid	SAMGrid operational support model enhancements
May-06	PPDG-OSG-Authorization	Implementation V1.1
May-06	PPDG-SRM	Multiple implementations of SRM V3
Jun-06	PPDG-OSG	Sustained exploitation of a significant fraction of the DOE-HENP resources
Jun-06	PPDG-OSG-Accounting	Implementation of Accounting V2
Jun-06	PPDG-SAMGrid	Analysis and production access to general HENP Grid resources for CDF and D0
Jun-06	PPDG-SAMGrid	Further develop accounting and authorization capabilities
Jun-06	PPDG-ATLAS	OSG deployment integrated with more general ATLAS analysis environment developed in collaboration with ARDA
Jun-06	PPDG-JLAB	Integrate VO tools into the production environment
Jun-06	PPDG-STAR	STAR will aim to achieve full access to Grid resources at universities and collaborating institutions
Jun-06	PPDG-OSG-Troubleshooting	Upgraded implementation
Jul-06	PPDG-OSG-Authorization	Implementation V2

5.3 Meetings

The following table shows the meetings held this quarter. The calendar at http://www.ppdg.net/calendars.htm has links to the agendas and minutes.

Oct 1	**12:30 p.m. - 1:30 p.m.**	**PDPG phone meeting - monitoring**
Oct 2	**10 a.m. - 11 a.m.**	**PPDG CS-11 phone conference**
Oct 10	**9 a.m. - 10 a.m.**	**PPDG Data Mgmt phone con**
Oct 16	**10 a.m. - 11 a.m.**	**PPDG CS-11 phone conference**

Oct 22	12:30 p.m. - 2 p.m.	PPDG weekly meeting
Oct 24	9 a.m. - 10 a.m.	PPDG Data Mgmt area meeting
Oct 29	12:30 p.m. - 2 p.m.	PPDG steering meeting
Oct 30	10 a.m. - 11 a.m.	PPDG CS-11 phone conference
Nov 5	12:30 p.m. - 2 p.m.	PPDG weekly meeting
Nov 7	All Day	PPDG face-to-face steering committee meeting, FNAL
Nov 12	12:30 p.m. - 2 p.m.	PPDG steering meeting
Dec 3	12:30 p.m. - 2 p.m.	PPDG weekly meeting
Dec 11	10 a.m. - 11 a.m.	PPDG CS-11 phone conference
Dec 15	All Day	PPDG Collaboration Meeting, LBNL, Berkeley
Dec 16	All Day	PPDG Collaboration Meeting, LBNL, Berkeley
Jan 7	12:30 p.m. - 2 p.m.	PPDG Steering Meeting
Jan 14	12:30 p.m. - 2 p.m.	PPDG weekly meeting
Jan 21	12:30 p.m. - 2 p.m.	PPDG, iVDGL (Trillium) - DOEGrids CA PKI teleconference
Jan 28	12:30 p.m. - 2 p.m.	PPDG weekly meeting
Feb 4	12:30 p.m. - 2 p.m.	PPDG weekly meeting
Feb 18	12:30 p.m. - 2 p.m.	PPDG weekly meeting
Feb 25	12:30 p.m. - 2 p.m.	PPDG Data Management meeting
Mar3	12:30 p.m. - 2 p.m.	PPDG weekly meeting